

Leveraging User Annotations in Sentiment Summarization

Ryan McDonald



Joint work with Ivan Titov (Geneva),
Sasha Blair-Goldensohn, Kerry Hannan,
Tyler Neylon, Jeff Reynar (Google)

User Generated (Text) Content

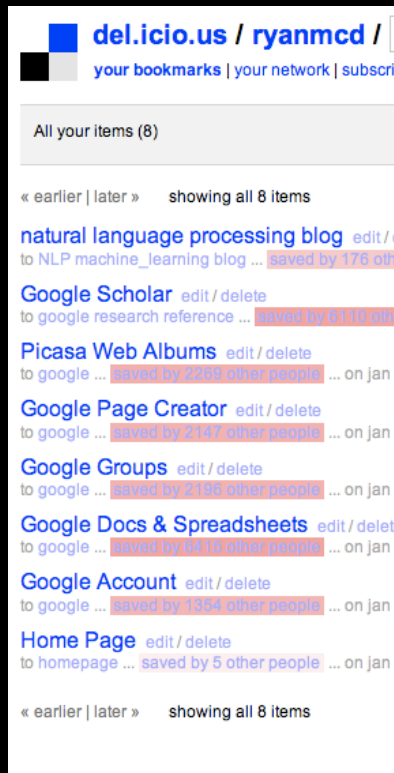
User Generated (Text) Content



User Generated (Text) Content



The screenshot shows the Wikipedia article for Abraham Lincoln. At the top left is the Wikipedia logo with the text "WIKIPEDIA The Free Encyclopedia". Below it is a navigation menu with items like "Main Page", "Contents", "Featured content", "Current events", and "Random article". The main content area has tabs for "article", "discussion", "view source", and "history". The article title is "Abraham Lincoln" and it starts with "From Wikipedia, the free encyclopedia". A note says "For other uses, see Abraham Lincoln (disambiguation)". The main text begins with "Abraham Lincoln (February 12, 1809 – April 15, 1865) was the sixteenth President of the United States..." and continues with details about his presidency and the American Civil War.



The screenshot shows a del.icio.us bookmark page for user ryanmcd. The header includes the del.icio.us logo and the user's name. Below the header, there are links for "your bookmarks", "your network", and "subscriptions". The main content area is titled "All your items (8)" and shows a list of 8 bookmarked items. Each item includes a title, a brief description, and a "saved by" count. The items listed are: "natural language processing blog", "Google Scholar", "Picasa Web Albums", "Google Page Creator", "Google Groups", "Google Docs & Spreadsheets", "Google Account", and "Home Page".

User Generated (Text) Content

The image shows a composite of two web pages. The top page is a Wikipedia article for "Cafe D'Alsace" in New York City, featuring a map and a user review from Ryan dated July 25, 2007. The bottom page is a del.icio.us bookmark page for the same user, showing a list of saved items including Google Scholar, Picasa Web Albums, Google Page Creator, Google Groups, Google Docs & Spreadsheets, and Google Account.

Wikipedia Article: Cafe D'Alsace

From Wikipedia, the free encyclopedia

Cafe D'Alsace
For other uses, see **Cafe D'Alsace** (disambiguation).
★★★★☆ 89 reviews

1695 2nd Ave
New York, NY 10128
(212) 722-5133
uppereast.com

Get Directions: [To here](#) - [From here](#)

[View Larger Map](#)

Your review
★★★★☆ **Good restaurant on UES** - Ryan - Jul 25, 2007
Cafe D'Alsace is one of the best restaurants I have found on the upper east side. The food is very good, even the simple things like a hangar steak and fries is cooked well and tasty. The cheese appetizer is excellent, though I wish they would occasionally change some of the options. Great beer and wine selection, beers mostly from bottle. Though they say you should make reservations, if there is just two of you, they will almost certainly squeeze you in without too much of a wait. Any more and I would definitely make sure you reserve.
[Edit](#) - [Delete](#)

All reviews
★★★★☆ **Cafe d'Alsace** - May 24, 2006
Will Upper East Siders get their fill of choucroute garni and baeckoffe in the neighborhood formerly home to a bevy of German restaurants? Strasbourg meets Schaller und Weber? ...
Was this review helpful? [Yes](#) - [No](#)
[More from NewYorkCity.com »](#)

★★★★☆ **Café D'Alsace Restaurant New York New...**
Upper East Siders, starved for well-prepared food in a sleek setting, had often migrated downtown when their stomachs rumbled. But **Café D'Alsace** gives locals one less reason to ...
Was this review helpful? [Yes](#) - [No](#)
[More from Gayot.com »](#)

del.icio.us / ryanm
your bookmarks | your network |

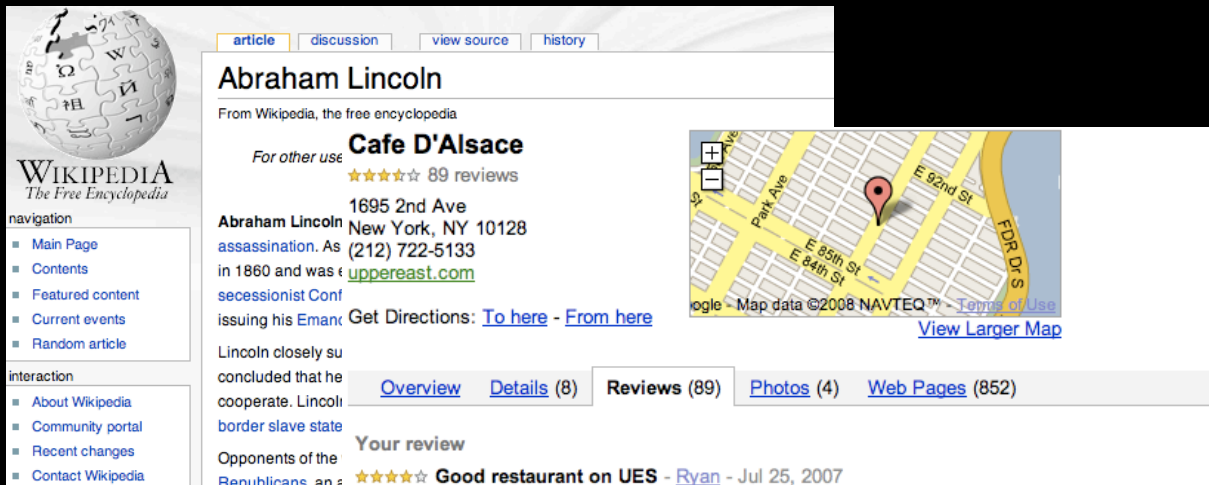
All your items (8)

« earlier | later » showing all 8 items

- [natural language processing blog](#) to NLP machine_learning blog ... [saved by](#)
- [Google Scholar](#) edit / delete to google research reference ... [saved by](#)
- [Picasa Web Albums](#) edit / delete to google ... [saved by 2269 other people](#) ...
- [Google Page Creator](#) edit / delete to google ... [saved by 2147 other people](#) ...
- [Google Groups](#) edit / delete to google ... [saved by 2196 other people](#) ... on Jan
- [Google Docs & Spreadsheets](#) edit / delete to google ... [saved by 6419 other people](#) ... on Jan
- [Google Account](#) edit / delete to google ... [saved by 1364 other people](#) ... on Jan
- [Home Page](#) edit / delete to homepage ... [saved by 5 other people](#) ... on Jan

« earlier | later » showing all 8 items

User Generated (Text) Content



WIKIPEDIA The Free Encyclopedia

article | discussion | view source | history

Abraham Lincoln

From Wikipedia, the free encyclopedia

Cafe D'Alsace

For other uses, see **Cafe D'Alsace** (disambiguation).

1695 2nd Ave
New York, NY 10128
(212) 722-5133
uppereast.com

Get Directions: [To here](#) - [From here](#)

View Larger Map

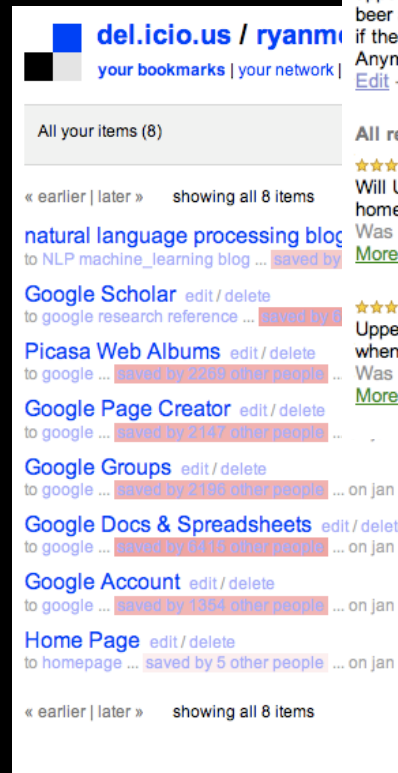
Overview | Details (8) | Reviews (89) | Photos (4) | Web Pages (852)

Your review

★★★★☆ **Good restaurant on UES** - Ryan - Jul 25, 2007

Cafe D'Alsace is one of the best restaurants I have found on the upper east side. The food is very good, even the simple things like a hangar steak and fries is cooked well and tasty. The cheese appetizer is excellent, though I wish they would occasionally change some of the options. Great beer and wine selection, beers mostly from bottle. Though they say you should make reservations, if there is just two of you, they will almost certainly squeeze you in without too much of a wait. Anymore and I would definitely make sure you reserve.

[Edit](#) - [Delete](#)



del.icio.us / ryanm

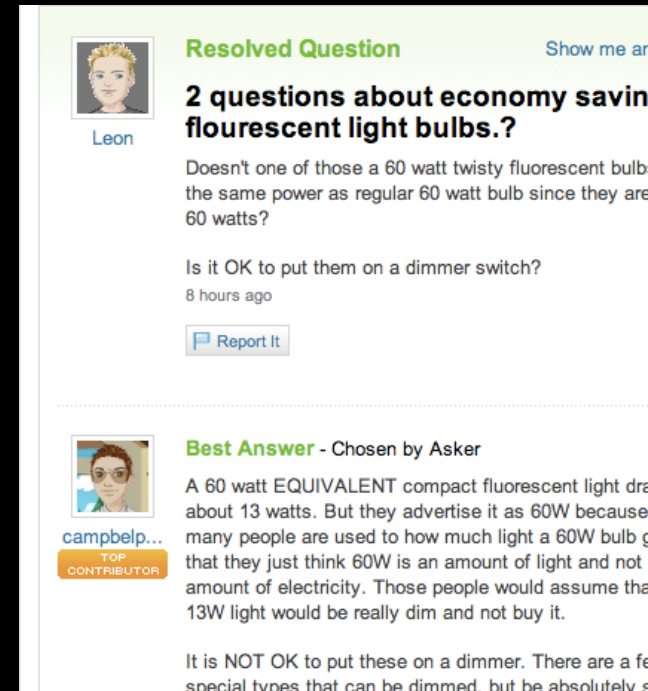
your bookmarks | your network |

All your items (8)

« earlier | later » showing all 8 items

- [natural language processing blog](#) to NLP machine_learning blog ... saved by [ryanm](#)
- [Google Scholar](#) edit / delete to google research reference ... saved by [ryanm](#)
- [Picasa Web Albums](#) edit / delete to google ... saved by 2209 other people ...
- [Google Page Creator](#) edit / delete to google ... saved by 2147 other people ...
- [Google Groups](#) edit / delete to google ... saved by 2196 other people ... on Jan 2007
- [Google Docs & Spreadsheets](#) edit / delete to google ... saved by 6419 other people ... on Jan 2007
- [Google Account](#) edit / delete to google ... saved by 1364 other people ... on Jan 2007
- [Home Page](#) edit / delete to homepage ... saved by 5 other people ... on Jan 2007

« earlier | later » showing all 8 items



Resolved Question

Show me an answer

2 questions about economy saving fluorescent light bulbs.?

Leon

Doesn't one of those a 60 watt twisty fluorescent bulb use about the same power as regular 60 watt bulb since they are 60 watts?

Is it OK to put them on a dimmer switch?
8 hours ago

[Report It](#)

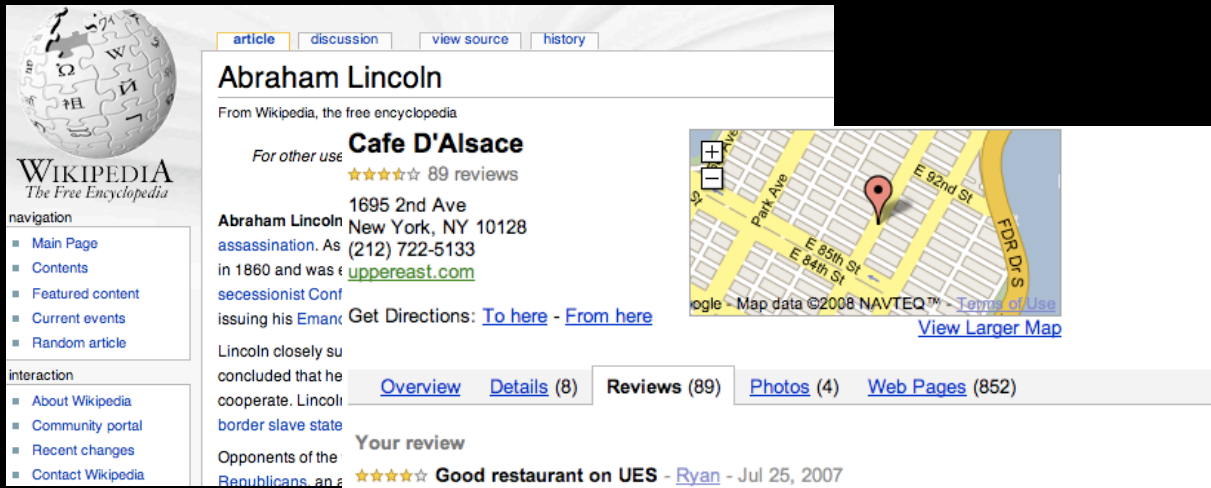
Best Answer - Chosen by Asker

[cambelp...](#)
TOP CONTRIBUTOR

A 60 watt EQUIVALENT compact fluorescent light draws about 13 watts. But they advertise it as 60W because many people are used to how much light a 60W bulb gives that they just think 60W is an amount of light and not an amount of electricity. Those people would assume that a 13W light would be really dim and not buy it.

It is NOT OK to put these on a dimmer. There are a few special types that can be dimmed, but be absolutely sure

User Generated (Text) Content



WIKIPEDIA The Free Encyclopedia

Abraham Lincoln

From Wikipedia, the free encyclopedia

Cafe D'Alsace

For other uses, see **Cafe D'Alsace** (disambiguation).

★★★★☆ 89 reviews

1695 2nd Ave
New York, NY 10128
(212) 722-5133
uppereast.com

Get Directions: [To here](#) - [From here](#)

[View Larger Map](#)

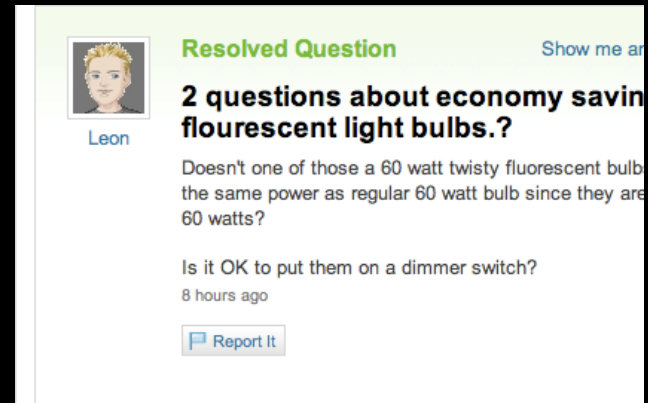
Overview Details (8) Reviews (89) Photos (4) Web Pages (852)

Your review

★★★★☆ **Good restaurant on UES** - Ryan - Jul 25, 2007

Cafe D'Alsace is one of the best restaurants I have found on the upper east side. The food is very good, even the simple things like a hangar steak and fries is cooked well and tasty. The cheese appetizer is excellent, though I wish they would occasionally change some of the options. Great beer and wine selection, beers mostly from bottle. Though they say you should make reservations, if there is just two of you, they will almost certainly squeeze you in without too much of a wait. Anymore and I would definitely make sure you reserve.

[Edit](#) - [Delete](#)



Resolved Question [Show me all](#)

Leon

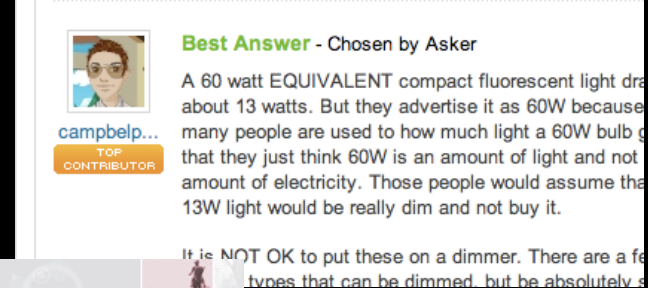
2 questions about economy saving fluorescent light bulbs.?

Doesn't one of those a 60 watt twisty fluorescent bulb use the same power as regular 60 watt bulb since they are 60 watts?

Is it OK to put them on a dimmer switch?

8 hours ago

[Report It](#)

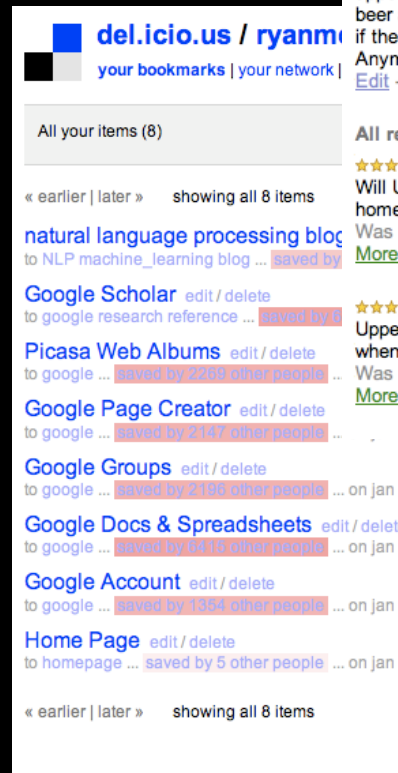


campbelp... **TOP CONTRIBUTOR**

Best Answer - Chosen by Asker

A 60 watt EQUIVALENT compact fluorescent light draws about 13 watts. But they advertise it as 60W because many people are used to how much light a 60W bulb gives that they just think 60W is an amount of light and not amount of electricity. Those people would assume that a 13W light would be really dim and not buy it.

It is NOT OK to put these on a dimmer. There are a few types that can be dimmed, but be absolutely sure



del.icio.us / ryanm

your bookmarks | your network

All your items (8)

« earlier | later » showing all 8 items

[natural language processing blog](#) to NLP machine_learning blog ... saved by [Google Scholar](#) edit / delete to google research reference ... saved by [Picasa Web Albums](#) edit / delete to google ... saved by 2269 other people ... [Google Page Creator](#) edit / delete to google ... saved by 2147 other people ... [Google Groups](#) edit / delete to google ... saved by 2196 other people ... on Jan ... [Google Docs & Spreadsheets](#) edit / delete to google ... saved by 6419 other people ... on Jan ... [Google Account](#) edit / delete to google ... saved by 1364 other people ... on Jan ... [Home Page](#) edit / delete to homepage ... saved by 5 other people ... on Jan

« earlier | later » showing all 8 items



natural language processing blog

03 February 2008

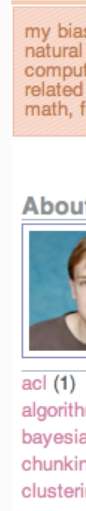
★★★★☆ **Cafe d'Alsace**
Will Upper East Siders (home to a bevy of German) Was this review helpful? [More from NewYorkCity](#)

★★★★☆ **Café D'Alsace**
Upper East Siders, start when their stomachs rumble. Was this review helpful? [More from Gayot.com](#)

The behemoth, PubMed

The friend I crashed with while attending SODA is someone I've known since we were five years old. (Incidentally, there's actually someone in the NLP world who I've actually known from earlier...small world.) Anyway, the friend I stayed with is just finishing med school at UCSF and will soon be staying there for residency. His specialty is neurosurgery, and his interests are in neural pathologies. He spent some time doing research on Alzheimer's disease, effectively by studying mice (there's something I feel sort of bad about finding slightly amusing about mice with Alzheimer's disease). Needless to say, in the process of doing research, he made nearly daily use out of PubMed. (For those of you who don't know, PubMed is like the ACL anthology, but with hundreds of thousands of papers, with new ones being added by the truckload daily, and will a bunch of additional things, like ontologies and data sets.)

There are two things I want to talk about regarding PubMed. I think both of these admit very interesting problems that we, as NLPers, are qualified to tackle. I think the most important



my bias
natural
comput
related
math, fu

About

acl (1)
algorithm
bayesian
chunking
clustering

User Generated (Text) Content

Wikipedia article for Cafe D'Alsace. The article includes the name, address (1695 2nd Ave, New York, NY 10128), phone number (212) 722-5133, and website (uppereast.com). It also features a map of the location and a 'Your review' section with a 5-star rating and text: "Cafe D'Alsace is one of the best restaurants I have found on the upper east side. The food is very good, even the simple things like a hangar steak and fries is cooked well and tasty. The cheese appetizer is excellent, though I wish they would occasionally change some of the options. Great beer and wine selection, beers mostly from bottle. Though they say you should make reservations, if there is just two of you, they will almost certainly squeeze you in without too much of a wait. Anymore and I would definitely make sure you reserve." Navigation tabs for Overview, Details, Reviews, Photos, and Web Pages are visible.

del.icio.us / ryanm... your bookmarks | your network | All your items (8) showing all 8 items. The list includes: natural language processing blog, Google Scholar, Picasa Web Albums, Google Page Creator, Google Groups, Google Docs & Spreadsheets, Google Account, and Home Page.

natural language processing blog. The main article is titled "The behemoth, PubMed" and discusses the challenges of managing a large volume of research papers. The text includes: "The friend I crashed with while attending SODA is someone I've known since we were five years old. (Incidentally, there's actually someone in the NLP world who I've actually known from earlier...small world.) Anyway, the friend I stayed with is just finishing med school at UCSF and will soon be staying there for residency. His specialty is neurosurgery, and his interests are in neural pathologies. He spent some time doing research on Alzheimer's disease, effectively by studying mice (there's something I feel sort of bad about finding slightly amusing about mice with Alzheimer's disease). Needless to say, in the process of doing research, he made nearly daily use out of PubMed. (For those of you who don't know, PubMed is like the ACL anthology, but with hundreds of thousands of papers, with new ones being added by the truckload daily, and will a bunch of additional things, like ontologies and data sets.) There are two things I want to talk about regarding PubMed. I think both of these admit very interesting problems that we, as NLPers, are qualified to tackle. I think the most important

Resolved Question: 2 questions about economy saving fluorescent light bulbs.? by Leon. The thread includes a question about whether 60-watt twisty fluorescent bulbs are the same power as regular 60-watt bulbs and if they can be used on a dimmer switch. A reply from sunya provides information on customizing Google pages. The forum interface shows user avatars, profile links, and interaction options like 'Report It' and 'Reply to author'.

User Generated Content

- Abundant source of information
- Diverse source (expert to novice to biased to spam)
- Data often contains structured labels

Now, don't get me wrong -- I'm not saying that this is easy -- but for some people who are constantly looking for "natural artifacts" of their craft, this is an enormous task.

Posted by hal at 2/03/2008 08:05:00 AM |

Labels: [problems](#), [summarization](#)

★★★★☆ **Good restaurant on UES**
Cafe D'Alsace is one of the best restaurants in the city. The food is good, even the simple things like an appetizer is excellent, though I wish

[del.icio.us](#) / tag / [nlp](#)
your bookmarks | your network | subscribe

All items tagged [nlp](#) → view [yours](#), [popular](#)

« [earlier](#) | [later](#) »

[The Stanford NLP \(Natural Language Processing\) class](#)
by [jasonzou273](#) to [NLP chinese-segmentation](#) ... 31 minutes ago

[TnT -- Statistical Part-of-Speech Tagging](#)
appears to be a decade old
by [onedanshow](#) to [research nlp](#) ... [save](#)

[CS 224N / Ling 280 Syllabus](#)
stanford class on natural language processing
by [neveaire](#) to [nlp](#) ... [saved by 3 other people](#) ... 1

Awesome!

[✎](#) Reviewed By: Mel on 11/11/08

[📄](#) Tech Level: somewhat high

Pros: This camera has tons of different picture taking scenes. You can take panoramic pictures of a scene. It's very handy.

Cons: The only con is of course many problems with it because of its age. However, I have found an app that you can charge it from an iPod.

Most Helpful Customer Reviews

291 of 310 people found the following review helpful:

★★★★☆ [Apple improves the iPhone and delivers an](#)

What to do with annotations?

- Traditional view: use them as training data
 - Train a model and run/eval it on new data
 - Isn't this just a contrived task?
 - Not always:
 - Train sentiment classifiers on reviews use it 4 blogs
 - Train review ratings, apply on phrases or sentences
 - Train on one blog, apply to unannotated blogs

Leverage Annotations for Related Problems

- Can we use ...?
 - Star ratings to predict phrase level sentiment
 - Star ratings to segment the text
 - Del.icio.us tags to place ads / improve ranking
 - Helpfulness rankings to extract QA pairs
 - Helpfulness rankings to build language models
 - ...
- Auxiliary tasks are closely related to signals provided by the user

Text Segmentation

- Focus on models for segmenting text
 - Can we use aspect ratings?

Food: 5; Decor: 5; Service: 5; Value: 5


The chicken was great. On top of that our service was excellent and the price was right. Can't wait to go back!

Food: 2; Decor: 1; Service: 3; Value: 2

We went there for our anniversary. My soup was cold and expensive plus it felt like they hadn't painted since 1980.

Food: 3; Decor: 5; Service: 4; Value: 5

The food is only mediocre, but well worth the cost. Wait staff was friendly. Lot's of fun decorations.



Food

"The chicken was great", "My soup was cold", "The food is only mediocre"

Decor

"it felt like they hadn't painted since 1980", "Lots of fun decorations"

Service

"service was excellent",
"Wait staff was friendly"

Value

"the price was right", "My soup was cold and expensive", "well worth the cost"

Sentiment Summarization

- Take a set of reviews for an entity and summarize them
- **Aspect-based summarization** (Hu & Liu 2004)
 - Summarize along key aspects

Nikos' Fine Dining

Food	4/5	“Best fish in the city”, “Excellent appetizers”
Decor	3/5	“Cozy with an old world feel”, “Too dark”
Service	1/5	“Our waitress was rude”, “Awful service”
Value	5/5	“Good Greek food for the \$”, “Great price!”

- Many real world manual examples, e.g., Zagat.com

Previous Work

- Hu and Liu '04
 - Aspect-based summarization
 - String-based aspects + lexicon sentiment
- Popescu and Etzioni '05: Opine system
- Gammon et al. '05
 - Aspect clusters: use most frequent word label
- Carenini '06
 - String-based + ontologies
- Mei et al. '07
 - Generative topic-sentiment models (at document level)

Three Tasks

Nikos' Fine Dining

Food	4/5	“Best fish in the city”, “Excellent appetizers”
Decor	3/5	“Cozy with an old world feel”, “Too dark”
Service	1/5	“Our waitress was rude”, “Awful service”
Value	5/5	“Good Greek food for the \$”, “Great price!”

Three Tasks

- Identify Aspects
 - Often we know this (pros-cons, tech specs, ontologies)

Nikos' Fine Dining

Food	4/5	“Best fish in the city”, “Excellent appetizers”
Decor	3/5	“Cozy with an old world feel”, “Too dark”
Service	1/5	“Our waitress was rude”, “Awful service”
Value	5/5	“Good Greek food for the \$”, “Great price!”

Three Tasks

- Identify Aspects
 - Often we know this (pros-cons, tech specs, ontologies)
- Extract Mentions
 - We always have to do this

Nikos' Fine Dining

Food	4/5	“Best fish in the city”, “Excellent appetizers”
Decor	3/5	“Cozy with an old world feel”, “Too dark”
Service	1/5	“Our waitress was rude”, “Awful service”
Value	5/5	“Good Greek food for the \$”, “Great price!”

Three Tasks

- Identify Aspects
 - Often we know this (pros-cons, tech specs, ontologies)
- Extract Mentions
 - We always have to do this
- Aggregate Sentiment
 - Again, we often know this (star ratings, eg, TripAdvisor)

Nikos' Fine Dining

Food	4/5	“Best fish in the city”, “Excellent appetizers”
Decor	3/5	“Cozy with an old world feel”, “Too dark”
Service	1/5	“Our waitress was rude”, “Awful service”
Value	5/5	“Good Greek food for the \$”, “Great price!”

Three Tasks

- Identify Aspects
 - Often we know this (pros-cons, tech specs, ontologies)
- Extract Mentions
 - We always have to do this
- Aggregate Sentiment
 - Again, we often know this (star ratings, eg, TripAdvisor)

Nikos' Fine Dining

Food	4/5	“Best fish in the city”, “Excellent appetizers”
Decor	3/5	“Cozy with an old world feel”, “Too dark”
Service	1/5	“Our waitress was rude”, “Awful service”
Value	5/5	“Good Greek food for the \$”, “Great price!”

Aspect Identification and Extraction

Nikos' Fine Dining

Food	4/5	"Best fish in the city", "Excellent appetizers"
Decor	3/5	"Cozy with an old world feel", "Too dark"
Service	1/5	"Our waitress was rude", "Awful service"
Value	5/5	"Good Greek food for the \$", "Great price!"

- Common method: **String-based extraction**
 - Find frequently occurring nouns that are modified by opinion words
 - Take top K as relevant aspects
 - Extract all sentences / phrases that match
 - **Problem:** Get a long list of aspects w/ no clustering

Aspect Identification and Extraction

- String-based example: restaurants
 - Is list really summarization?
 - How far down to get “cozy”, “fish”, “\$”, “waitress”, “dark”?
- We really want to cluster these

General
Food
Ambiance
Service
Value

food
place
pizza
service
restaurant
atmosphere
time
wine
meal
prices
value
sauce
hour
price
dim
selection
experience
crust
dining
ingredients

Nikos' Fine Dining


Food	4/5	“Best fish in the city”, “Excellent appetizers”
Decor	3/5	“Cozy with an old world feel”, “Too dark”
Service	1/5	“Our waitress was rude”, “Awful service”
Value	5/5	“Good Greek food for the \$”, “Great price!”

Topic Models

- Studied in ML and Data Mining
 - LSA, PLSA, LDA, Pachinko Allocation, ...
- Build semantic “topics” of data collections
 - e.g., newsgroups into “religion”, “politics”, “science”, ...
- **Simple hypothesis**
 - Topics in reviews correspond to clustered aspects
- We will focus on LDA type models (Blei et al. '03)
 - Others produce similar observations

LDA

- Generative model of text
- Sample multinomial word distributions for each topic
- The for each document d :



- choose distribution of topics $\theta_d \sim Dir(\alpha)$
- for each word i in document d
 - choose topic $z_{d,i} \sim \theta_d$,
 - choose word $w_{d,i} \sim \varphi_{z_{d,i}}$.

Side Note: Inference

- All methods use collapsed Gibbs (Griffiths & Steyvers '04)
- A sample from the chain used to approx:
 - Distribution of words in topics
 - Distribution of topics in text fragments
- We tried variational techniques, but they didn't work
- Not going to go into details
 - See Titov & McDonald (WWW 2008) for more

LDA

- Problem with LDA (and most other topic models)
 - Co-occurrences modeled at document level
 - Topics are about instances not aspects
 - e.g., iPod versus Creative Labs
 - Often clusters are meaningless

(Service??) Topic 0: product player did support bought work unit problem \$
(Creative Labs) Topic 1: gigabeat deleted waiting jukebox creative playback
(iPod) Topic 11: ipod apple mac firewire dock itunes x display aac

Most topics are incoherent. Only 4 out of first 40 can be viewed as aspects.

LDA

- Simple solutions: LDA over sentences
 - Co-occurrence counts too sparse
 - Can use sliding window, but results look like LDA
 - Still can't distinguish aspect topics from the rest
- Another solution: **Multi-grain topic models**
 - Model **local topics** (aspects) and **global topics** (types)
 - Creates a **bottleneck** for local topics
 - Words generated from sliding window

MG-LDA Topics (Mp3)

First 8 Local Topics!!

Sound Quality	Features	PC Connection	Tech Problems	Looks	Controls	Battery	Accessor's
sound	games	usb	reset	case	button	battery	usb
quality	features	pc	noise	pocket	play	hours	cable
headphones	clock	windows	backlight	silver	track	life	headphones
volume	contacts	port	slow	screen	menu	batteries	adapter
bass	calendar	transfer	freeze	plastic	song	charge	remote
earphones	alarm	computer	turn	clip	buttons	aaa	plug
ear	notes	mac	remove	easily	volume	rechargeable	power
rock	game	software	playing	small	album	time	charger
settings	quiz	cable	hot	blue	tracks	power	included

First 4 Global Topics

iPod	Creative Zen	Sony Walkman	Video Players
ipod	zen	sony	video
music	creative Zen	walkman	screen
apple	micro	memory	videos
songs	touch	stick	device
use	xtra	sonicstage	photos
mini	pad	players	tv
very	nomad	atrac3	archos
just	waiting	mb	pictures
itunes	labs	atrac	camera

LDA

- 40 topics
- Only 4 aspect topics
- A couple other coherent topics
- Good topics in no order
- Mostly junk topics

MG-LDA Topics (Mp3)

First 8 Local Topics!!

Sound Quality	Features	PC Connection	Tech Problems	Looks	Controls	Battery	Accessor's
sound quality headphones volume bass earphones ear rock settings	games features clock contacts calendar alarm notes	usb pc windows port transfer computer mac	reset noise backlight slow freeze turn remove	case pocket silver screen plastic clip easily	button play track menu song buttons volume	battery hours life batteries charge aaa rechargeable	usb cable headphones adapter remote plug power charger included

- Works for many other domains
- Works for both LDA and PLSA
- Empirically better than LDA on classification tasks
- Titov and McDonald 2008 at WWW

First 4

iPod			
ipod music apple songs use mini very just itunes	creative Zen micro touch xtra pad nomad waiting labs	walkman memory stick sonicstage players atrac3 mb atrac	screen videos device photos tv archos pictures camera

- 40 topics
- Only 4 aspect topics
- A couple other coherent topics
- Good topics in no order
- Mostly junk topics

MG-LDA Results

- Clearly the set of topics is better than standard models
- But, we don't know topic labels a priori
- **Solution**
 - Let user annotations guide us
- Many things at our disposal
 - Tech specs
 - Pros-cons lists
 - **Aspect Ratings**

?	?	?
sound	games	usb
quality	features	pc
headphones	clock	windows
volume	contacts	port
bass	calendar	transfer
earphones	alarm	computer
ear	notes	mac
rock	game	software
settings	quiz	cable

Aspect Ratings

- Available on an increasing number of websites
- Give us two things
 - Important aspects
 - Signals that are correlated to the text

Aunt Chilada's ☪ NEW

3931 Whitney Avenue
Hamden, CT 06518-1519
Phone: (203) 230-4640
[See map](#)

Category: Mexican

Number of reviews to date: 4

[Ed Scargo \(02/01/2008\)](#)
[Tom Williams \(06/30/2007\)](#)
[Loris Feldman \(10/17/2006\)](#)
[Tom W Williams \(06/19/2004\)](#)

[Add your review!](#)
[Email a friend](#)

Ratings average:

Food ☺☺☺☺

Service ☺☺☺☺

Price/Value ☺☺☺☺☺

Atmosphere ☺☺☺☺☺

Overall ☺☺☺☺

This is a good place to get a drink but the food is not Tex-Mex (or Mexican). I would say it pretends to be Tex-Mex. Don't bother if even fairly decent Tex-Mex is what you are looking for. The food is boring, tasteless and unauthentic. The place can be loud. Service was mediocre. The waitress was inattentive. There is no decent wine list.

Submitted by:
Ed Scargo (02/01/2008)
Posted [2](#) reviews

[Printer-friendly review](#)

Food ☺☺

Service ☺☺

Price/Value ☺☺☺

Atmosphere ☺☺☺

Overall ☺☺

Was it a pleasant experience? No

How many were seated at your table? 2

Do they accept reservations? No

Would you return? No

If you ordered wine, what type was it? red

Credit cards accepted? Yes

[back to top](#)

We had a great time. Sat on the deck watching the golfers at the driving range. The best Margritas and Sasla I ever had. We had burritos that were very good also. The Guacamole could use a little

Food ☺☺☺☺☺

Service ☺☺☺☺☺

Price/Value ☺☺☺☺☺

RATING

☺ x5 - Excellent
☺☺ x4 - Average
☺☺☺ x3 - Average
☺☺☺☺ x2 - Below Average
☺☺☺☺☺ x1 - Poor

Get it on WeRthere!

Aspect Ratings

- Available on an increasing number of websites
- Give us two things
 - Important aspects
 - Significant correlations

Idea

Correlate topics with user provided aspect rankings
Topics are labeled!!
Topic quality should be much better!!

Aunt Chilada's

3931 Whitney Avenue
Hamden, CT 06518-1519
Phone: (203) 230-4640
[See map](#)

Category: Mexican

Number of reviews to date: 4

[Ed Scarco \(02/01/2008\)](#)
[Tom Williams \(06/30/2007\)](#)
[Loris Feldman \(10/17/2006\)](#)
[Tom W Williams \(06/19/2004\)](#)

Add your review!
Email a friend

Ratings average:	
Food	
Service	
Price/Value	
Atmosphere	
Overall	

This is a good place to get a drink but the food is not Tex-Mex (or Mexican). I would say it pretends to be Tex-Mex.
Don't bother if even fairly decent

Food	
Service	
Price/Value	

we had a great time. Sat on the deck watching the golfers at the driving range. The best Margritas and Sasla I ever had. We had burritos that were very good also. The Guacamole could use a little

Food	
Service	
Price/Value	

RATING

- x5 - Excellent
- x4 - Average
- x3 - Average
- x2 - Below Average
- x1 - Poor


Get your weather here!

Supervised LDA

- Can augment topic models to generate **observed signals**
 - S-LDA (Blei and McAuliffe '07 NIPS)
 - Use document labels to guide topic construction
- We take this insight and extend it to MG-LDA
 - For each aspect rating
 - Add a MaxEnt classifier to the model
 - Associate one topic to each classifier
 - MaxEnt classifier uses only words from that topic to predict rating

Multi-Aspect Sentiment Model

Can be any topic model




Multi-Aspect Sentiment Model

Can be any topic model

Overall sentiment variable

Models fact that aspect rankings are correlated

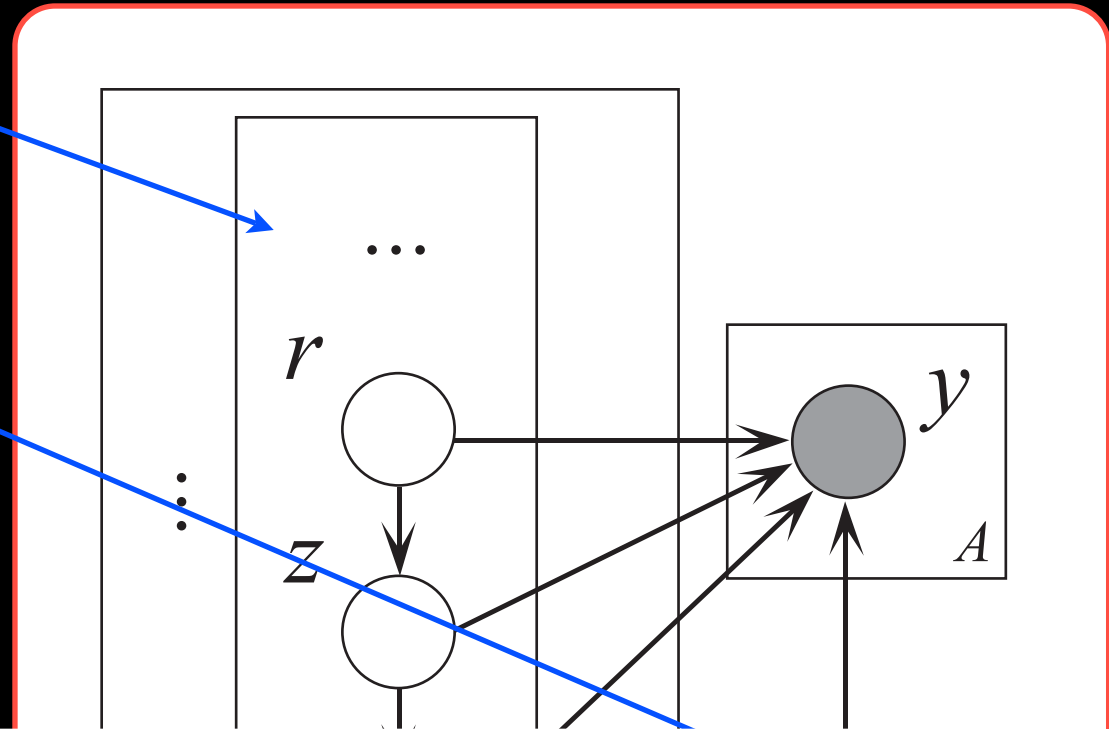


Multi-Aspect Sentiment Model

Can be any topic model

Overall sentiment variable

Models fact that aspect rankings are correlated



If we optimize the models jointly then topics will correspond directly to aspects

Multi-Aspect Sentiment Model

First 3 Local Topics

Service	Location	Rooms
staff	hotel	room
friendly	walk	bathroom
helpful	location	shower
service	station	bed
desk	metro	tv
concierge	walking	small
excellent	away	water
reception	right	clean
pleasant	minute	comfortable

- Hotel model with three aspects
 - service, location, rooms
 - Tied first three topics to these aspects ratings
 - Trained on 10,000 reviews
- Topics correspond to associated aspects!!

Multi-Aspect Sentiment Model

- Semantic models of aspects ... what does this buy us?
- Can use model directly to find mentions of aspects

Food: 5; Decor: 5; Service: 5; Value: 5


The chicken was great. On top of that our service was excellent and the price was right. Can't wait to go back!

Food: 2; Decor: 1; Service: 3; Value: 2

We went there for our anniversary. My soup was cold and expensive plus it felt like they hadn't painted since 1980.

Food: 3; Decor: 5; Service: 4; Value: 5

The food is only mediocre, but well worth the cost. Wait staff was friendly. Lot's of fun decorations.



Food

"The chicken was great", "My soup was cold", "The food is only mediocre"

Decor

"it felt like they hadn't painted since 1980", "Lots of fun decorations"

Service

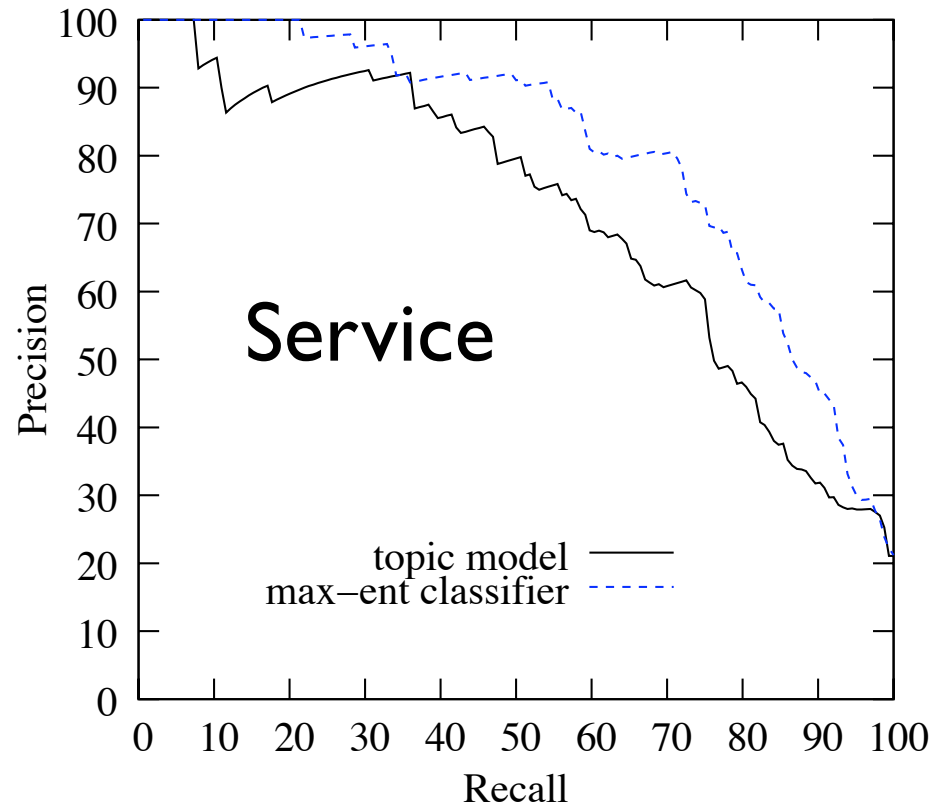
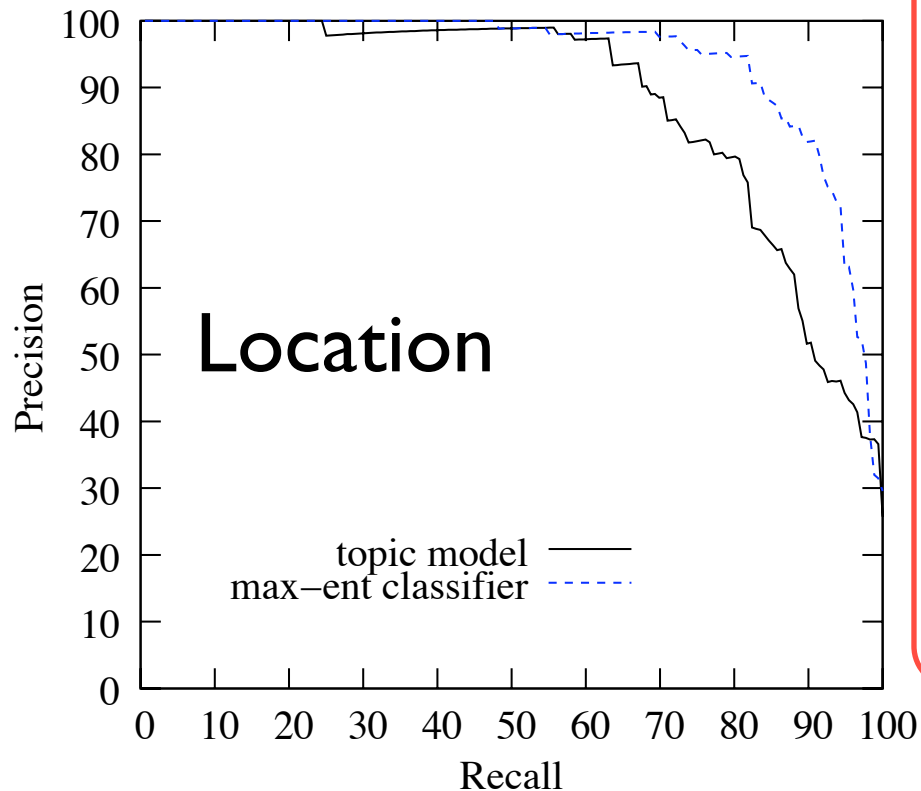
"service was excellent",
"Wait staff was friendly"

Value

"the price was right", "My soup was cold and expensive", "well worth the cost"

Aspect Mention Extraction

Compared to a supervised
MaxEnt model



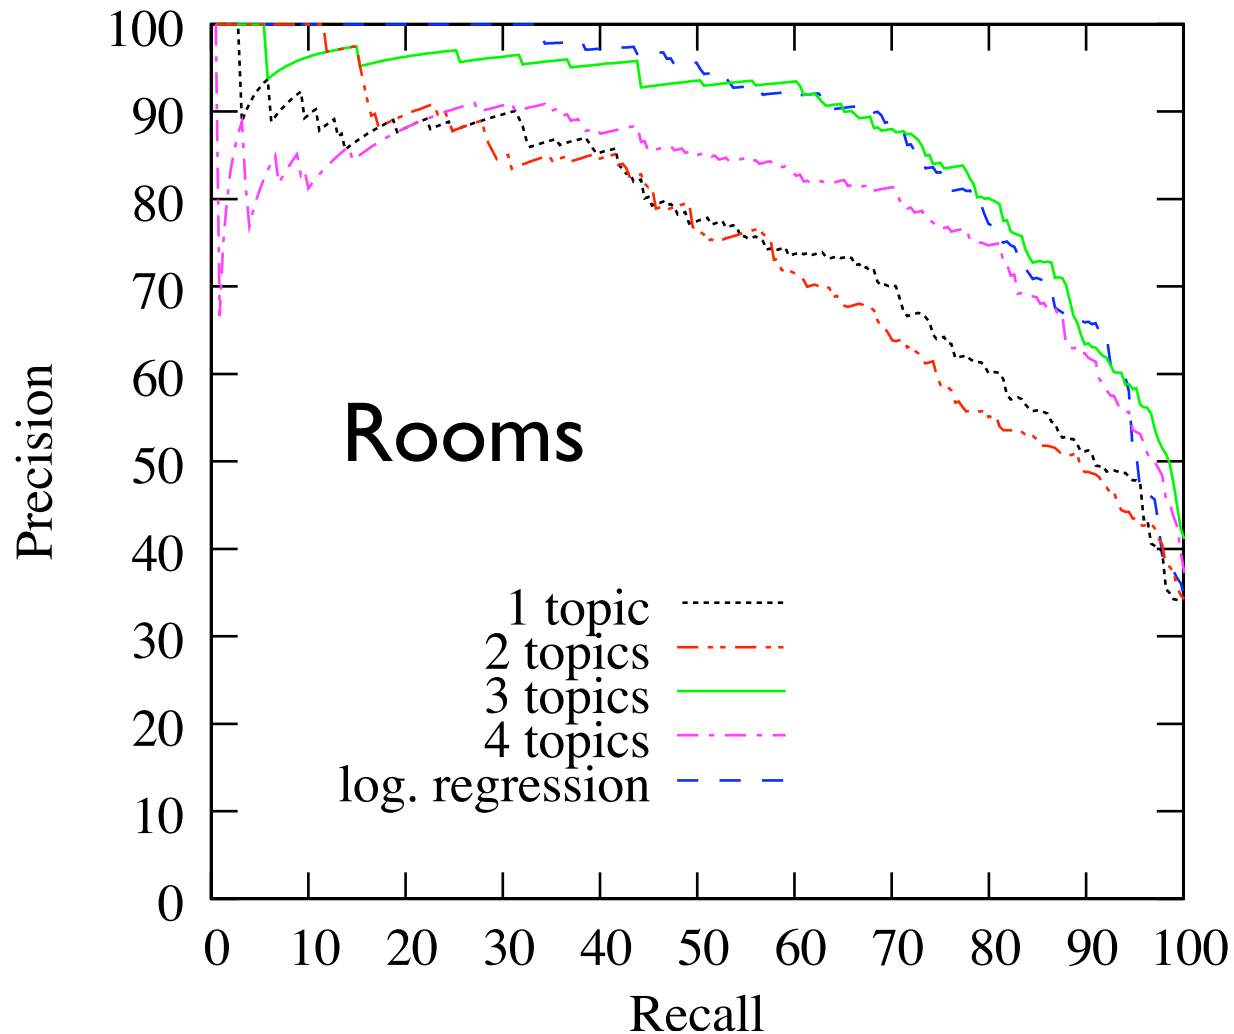
Multiple Topics per Aspect Classifier

- Required when an aspect is diverse
- e.g., Rooms = bed, bathroom, noise, view, ...

Rooms									
2		3			4				
rooms	room	room	room	room	room	check	room	bathroom	
clean	noise	clean	floor	bathroom	clean	arrived	noise	room	
hotel	night	bed	view	shower	rooms	time	night	shower	
room	street	comfortable	rooms	air	comfortable	day	night	tv	
small	did	rooms	suite	water	bed	airport	street	bed	
nice	air	bathroom	got	did	small	early	did	small	
comfortable	rooms	small	views	like	beds	room	air	water	
modern	door	beds	given	hot	nice	luggage	rooms	towels	
good	open	nice	quite	towel	bathroom	took	noisy	bath	

Aspect Mention Extraction

Multiple
topics
per
aspect



Topic Models & Sentiment Summ.

- Can topic models be used to summarize sentiment?
 - Yes!! For aspect identification and mention extraction
 - MG-LDA accurately finds aspect-like topics
 - But suffers from cluster labeling problem
 - Can augment MG-LDA to leverage aspect ratings
 - Ratings present in many data sets
 - Correlates topics w/ known aspects
 - Improves quality of topics
 - Yields highly precise mention extractors

Three Tasks

- Identify Aspects
 - Often we know this (pros-cons lists, star ratings)
- Extract Mentions
 - We always have to do this
- **Aggregate Sentiment**
 - We often know this (star ratings, eg, TripAdvisor)
 - But there is still a lot of data w/ out this

Could use
topic models

Nikos' Fine Dining

Food	4/5	“Best fish in the city”, “Excellent appetizers”
Decor	3/5	“Cozy with an old world feel”, “Too dark”
Service	1/5	“Our waitress was rude”, “Awful service”
Value	5/5	“Good Greek food for the \$”, “Great price!”

Aggregate Sentiment

- Simple
 - Extract mentions for each aspect
 - Average sentiment over each of them
- Problem
 - Current sentiment classifiers are either:
 - Domain specific
 - Low in accuracy

Domain Independent Classifiers

- Build weighted semantic graph -- $A = (a_{ij})$ -- from WordNet
 - Use synonyms, antonyms
 - Like Hu and Liu '04 and Kim and Hovy '06, but with optimization
- Use label propagation from seed sets of positive, negative and neutral words

Domain Independent Classifiers

- Build weighted semantic graph -- $A = (a_{ij})$ -- from WordNet
 - Use synonyms, antonyms
 - Like Hu and Liu '04 and Kim and Hovy '06, but with optimization
- Use label propagation from seed sets of positive, negative and neutral words

$$\mathbf{s}_i^0 = \begin{cases} +1 & \text{if } w_i \in P \\ -1 & \text{if } w_i \in N \\ 0 & \forall w_i \in \text{WordNet} - P \cup N \end{cases}$$

Domain Independent Classifiers

- Build weighted semantic graph -- $A = (a_{ij})$ -- from WordNet
 - Use synonyms, antonyms
 - Like Hu and Liu '04 and Kim and Hovy '06, but with optimization
- Use label propagation from seed sets of positive, negative and neutral words

$$s_i^0 = \begin{cases} +1 & \text{if } w_i \in P \\ -1 & \text{if } w_i \in N \\ 0 & \forall w_i \in \text{WordNet} \end{cases}$$

$$a_{ij} = \begin{cases} 1 + \lambda & \text{if } i == j \\ +\lambda & \text{if } w_i \in \text{syn}(w_j) \ \& \ w_i \notin M \\ -\lambda & \text{if } w_i \in \text{ant}(w_j) \ \& \ w_i \notin M \\ 0 & \text{otherwise.} \end{cases}$$

Domain Independent Classifiers

- Build weighted semantic graph -- $A = (a_{ij})$ -- from WordNet
 - Use synonyms, antonyms
 - Like Hu and Liu '04 and Kim and Hovy '06, but with optimization
- Use label propagation from seed sets of positive, negative and neutral words

$$\mathbf{s}_i^0 = \begin{cases} +1 & \text{if } w_i \in P \\ -1 & \text{if } w_i \in N \\ 0 & \forall w_i \in \text{WordNet} - P \cup N \end{cases}$$

$$a_{ij} = \begin{cases} 1 + \lambda & \text{if } i == j \\ +\lambda & \text{if } w_i \in \text{syn}(w_j) \ \& \ w_i \notin M \\ -\lambda & \text{if } w_i \in \text{ant}(w_j) \ \& \ w_i \notin M \\ 0 & \text{otherwise.} \end{cases}$$

for $m := 1$ to M
 $\mathbf{s}^m := \mathbf{A} \mathbf{s}^{m-1}$

Domain Independent Classifiers

- Build weighted semantic graph -- $A = (a_{ij})$ -- from WordNet
- Use synonyms, antonyms
- Like Hu and Liu '04 and Kim and Hovy '06, but with optimization
- Use label propagation from seed sets of positive, negative and neutral words

$$\mathbf{s}_i^0 = \begin{cases} +1 & \text{if } w_i \in P \\ -1 & \text{if } w_i \in N \\ 0 & \forall w_i \in \text{WordNet} \end{cases}$$

$$a_{ij} = \begin{cases} 1 + \lambda & \text{if } i == j \\ +\lambda & \text{if } w_i \in \text{syn}(w_j) \ \& \ w_i \notin M \\ -\lambda & \text{if } w_i \in \text{ant}(w_j) \ \& \ w_i \notin M \\ 0 & \text{otherwise.} \end{cases}$$

$$\text{for } m := 1 \text{ to } M \\ \mathbf{s}^m := \mathbf{A} \mathbf{s}^{m-1}$$


$$\text{raw-score}(x) := \sum_{i=1}^n \mathbf{s}_i.$$

Meta Classifier

- Collect scores for:
 - Sentence / phrase
 - Previous & next sentence / phrase
 - Document
- Train a classifier on a labeled set of sentences/phrases
 - Use scores as features


Meta Classifier

- Collect scores for:
 - Sentence / phrase
 - Previous & next sentence / phrase
 - Document
- Train a classifier on a labeled set of sentences/phrases
 - Use scores as features



Meta Classifier

- Collect scores for:
 - Sentence / phrase
 - Previous & next sentence / phrase
 - Document
- Train a classifier on a labeled set of sentences/phrases
 - Use scores as features



Meta Classifier

- Collect scores for:
 - Sentence / phrase
 - Previous & next sentence / phrase
 - Document
- Train a classifier on a labeled set of sentences/phrases
 - Use scores as features

User reviews usually have ****overall**** sentiment


Overall sentiment highly correlated w/ phrase & sentence sentiment



Meta Classifier

- Collect scores for:
 - Sentence / phrase
 - Previous & next sentence / phrase
 - Document
- Train a classifier on a labeled set of sentences/phrases
 - Use scores as features

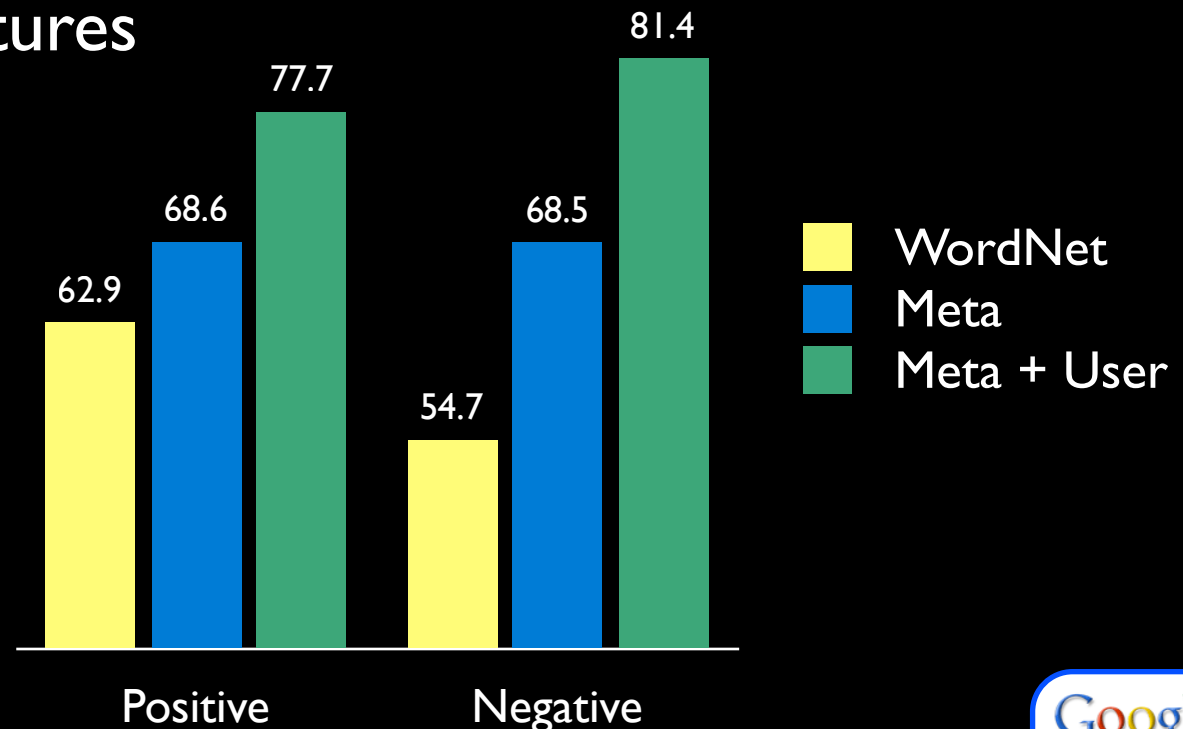
Add feature for overall user rating



Meta Classifier

- Collect scores for:
 - Sentence / phrase
 - Previous & next sentence / phrase
 - Document
- Train a classifier on a labeled set of sentences/phrases
 - Use scores as features

Add feature for overall user rating



Summary

- Aspect-based sentiment summarization
 - Can use topic models and standard sent classifiers
 - **Quality improves by leveraging correlated user signals**
 - Aspect ratings for aspect mention extraction
 - Overall ratings for phrase/sentence classification
- Models generalize to any segmentation problem where there are correlated signals
 - e.g., del.icio.us bookmarks, blog labels, helpfulness, ...

Thanks

- Joint work with Ivan Titov, Sasha Blair-Goldensohn, Kerry Hannan, Tyler Neylon, George Reis and Jeff Reynar
- Thank you to Kenji and Tsujii lab for invite